

More algorithms for provable dictionary learning

Sanjeev Arora* Aditya Bhaskara† Rong Ge‡ Tengyu Ma§

January 6, 2014

Abstract

In *dictionary learning*, also known as *sparse coding*, the algorithm is given samples of the form $y = Ax$ where $x \in \mathbb{R}^m$ is an unknown random sparse vector and A is an unknown dictionary matrix in $\mathbb{R}^{n \times m}$ (usually $m > n$, which is the *overcomplete* case). The goal is to learn A and x . This problem has been studied in neuroscience, machine learning, visions, and image processing. In practice it is solved by heuristic algorithms and provable algorithms seemed hard to find. Recently, provable algorithms were found that work if the unknown feature vector x is \sqrt{n} -sparse or even sparser. Spielman et al. [SWW12] did this for dictionaries where $m = n$; Arora et al. [AGM13] gave an algorithm for overcomplete ($m > n$) and incoherent matrices A ; and Agarwal et al. [AAN13] handled a similar case but with weaker guarantees.

This raised the problem of designing provable algorithms that allow sparsity $\gg \sqrt{n}$ in the hidden vector x . The current paper designs algorithms that allow sparsity up to $n/\text{poly}(\log n)$. It works for a class of matrices where features are *individually recoverable*, a new notion identified in this paper that may motivate further work.

The algorithm runs in quasipolynomial time because they use limited enumeration.

*Princeton University, Computer Science Department and Center for Computational Intractability. Email: arora@cs.princeton.edu. This work is supported by the NSF grants CCF-0832797, CCF-1117309, CCF-1302518, DMS-1317308, and Simons Investigator Grant.

†Google Research NYC. Email: bhaskara@cs.princeton.edu. Part of this work was done while the author was a Postdoc at EPFL, Switzerland.

‡Microsoft Research. Email: rongge@microsoft.com. Part of this work was done while the author was a graduate student at Princeton University and was supported in part by NSF grants CCF-0832797, CCF-1117309, CCF-1302518, DMS-1317308, and Simons Investigator Grant.

§Princeton University, Computer Science Department and Center for Computational Intractability. Email: tengyu@cs.princeton.edu. This work is supported by the NSF grants CCF-0832797, CCF-1117309, CCF-1302518, DMS-1317308, and Simons Investigator Grant.

1 Introduction

Dictionary learning, also known as *sparse coding*, tries to understand the structure of observed samples y by representing them as sparse linear combinations of “dictionary” elements. More precisely, there is an unknown dictionary matrix $A \in \mathbb{R}^{n \times m}$ (usually $m > n$, which is the *overcomplete* case), and the algorithm is given samples $y = Ax$ where x is an unknown random sparse vector. (We say a vector is k -sparse if it has at most k nonzero coordinates.) The goal is to learn A and x . Such sparse representation was first studied in neuroscience, where Olshausen and Field [OF97] suggested that dictionaries fitted to real-life images have similar properties as the receptive fields of neurons in the first layer of visual cortex. Inspired by this neural analog, dictionary learning is widely used in machine learning for *feature selection* [AEP06]. More recently the idea of sparse coding has also influenced deep learning [BC⁺07]. In image processing, learned dictionaries have been successfully applied to image denoising [EA06], edge detection [MLB⁺08] and super-resolution [YWHM08].

Provable guarantees for dictionary learning have seemed difficult because the obvious math programming formulation is nonconvex: both A and the x ’s are unknown. Even when the dictionary A is known, it is in general NP-hard to get the sparse combination x given *worst-case* y [DMA97]. This problem of decoding x given Ax with full knowledge of A is called *sparse recovery* or *sparse regression*, and is closely related to *compressed sensing*. For many types of dictionary A , sparse recovery was shown to be tractable even on worst-case y , starting with such a result for *incoherent* matrices by Donoho and Huo [DH01]. However in most early works x was constrained to be \sqrt{n} -sparse, until Candes, Romberg and Tao [CRT06] showed how to do sparse recovery even when the sparsity is $\Omega(n)$, assuming A satisfies the *restricted isometry property* (RIP) (which random matrices do).

But dictionary learning itself (recovering A given samples y) has proved much harder and heuristic algorithms are widely used. Lewicki and Sejnowski [LS00] designed the first one, which was followed by the method of optimal directions (MOD) [EAHH99] and K-SVD [AEB06]. See [Aha06] for more references. However, until recently there were no algorithms that provably recovers the correct dictionary. Recently Spielman et al. [SWW12] gave such an algorithm for the full rank case (i.e., $m = n$) and the unknown feature vector x is \sqrt{n} -sparse. However, in practice overcomplete dictionaries ($m > n$) are preferred. Arora et al. [AGM13] gave the first provable learning algorithm for overcomplete dictionaries that runs in polynomial-time; they required x to be $n^{1/2-\epsilon}$ -sparse (roughly speaking) and A to be incoherent. Independently, Agarwal et al. [AAN13] gave a weaker algorithm that also assumes A is incoherent and allows x to be $n^{1/4}$ -sparse. Thus all three of these recent algorithms cannot handle sparsity more than \sqrt{n} , and this is a fundamental limitation of the technique: they require two random x, x' to intersect in no more than $O(1)$ coordinates with high probability, which fails to hold when sparsity $\gg \sqrt{n}$. Since sparse recovery (where A is known) is possible even up to sparsity $\Omega(n)$, this raised the question whether dictionary learning is possible in that regime. In this paper we will refer to feature vectors with sparsity $n/\text{poly}(\log n)$ as *slightly-sparse*, since methods in this paper do not seem to allow density higher than that.

In our recent paper on deep learning ([ABGM13], Section 7) we showed how to solve dictionary learning in this regime for dictionaries which are adjacency matrices of random

weighted sparse bipartite graphs; these are known to allow sparse recovery albeit with a slight twist in the problem definition [Ind08, JXHC09, BGI⁺08]. Since real-life dictionaries are probably not random, this raises the question whether dictionary learning is possible in the slightly sparse case for other dictionaries. The current paper gives quasipolynomial-time algorithms for learning more such dictionaries. The running time is quasipolynomial time because it uses limited enumeration (similarly, e.g., to algorithms for learning gaussian mixtures). Now we discuss this class of dictionaries.

Some of our discussion below refers to nonnegative dictionary learning, which constrains matrices A and hidden vector x to have nonnegative entries. This is a popular variant proposed by Hoyer [Hoy02], motivated again partly by the neural analogy. Algorithms like NN-K-SVD [AEB05] were then applied to image classification tasks. This version is also related to *nonnegative matrix factorization* [LS99], which has been observed to lead to factorizations that are usually sparser and more local than traditional methods like SVD.

1.1 How to define dictionary learning?

Now we discuss what versions of dictionary learning make more sense than others. For exposition purposes we refer to the coordinates of the hidden vector x as *features*, and those of the visible vector $y = Ax$ as *pixels*, even though the discussion applies to more than just computer vision. Dictionary learning as defined here—which is the standard definition—assumes that features’ effect on the pixels add linearly.

But, the problem definition is somewhat arbitrary. On the one hand one could consider *more general* (and nonlinear) versions of this problem—for instance in vision, dictionary learning is part of a system that has to deal with *occlusions* among objects that may hide part of a feature, and to incorporate the fact that features may be present with an arbitrary translation/rotation. On the other hand, one could consider more specific versions that place restrictions on the dictionary, since not all dictionaries may make sense in applications. We consider this latter possibility now, with the usual caveat that it is nontrivial to cleanly formalize properties of real-life instances.

One reasonable property of real-life dictionaries is that each feature does not involve most pixels. This implies that column vectors of A are *relatively sparse*. Thus matrices with RIP property—at least if they are dense—do not seem a good match¹.

Another intuitive property is that features are *individually recoverable*, which means, roughly speaking, that to an observer who knows the dictionary, the presence of a particular feature should not be confusable with the effects produced by the usual distribution of other features (this is an average-case condition, since x satisfies stochastic assumptions). In particular, one should be able to detect its presence by looking only at the pixels it would affect.

Thus it becomes clear that not all matrices that allow sparse recovery are of equal interests. The paper of Arora et al. [AGM13] restricts attention to *incoherent* matrices, where the columns have pairwise inner product at most μ/\sqrt{n} where μ is small, like $\text{poly}(\log n)$. These make sense on both the above counts. First, they can have fairly sparse columns. Secondly they satisfy $A^T A \approx I$, so given Ax one can take its inner product with the i th

¹By contrast, in the usual setting of compressed sensing, the matrix provides a basis for making measurements, and its density is a nonissue.

column A_i to roughly determine the extent to which feature i is present. But incoherent matrices restrict sparsity to $O(\sqrt{n})$, so one is tempted by RIP matrices but, as mentioned, their columns are fairly dense. Furthermore, RIP matrices were designed to allow sparse recovery for worst-case feature vectors whereas in dictionary learning these are stochastic. As mentioned, sparse random graphs (with random edges weights in $[-1, 1]$) check all the right boxes (and were handled in our recent paper on deep learning) but require positing that the dictionary has no structure. The goal in the current paper is to move beyond random graphs.

Dictionaries with individually recoverable features. Let us try to formulate the property that features are *individually recoverable*. We hope this definition and discussion will stimulate further work (similar, we hope, to Dasgupta’s formalization of separability for gaussian mixtures [Das99]). Let us assume that the coordinates of x are pairwise independent. Then the presence of the i th feature (i.e., $x_i \neq 0$) changes the conditional distribution of those pixels involved in A_i , the i th column of A . Features are said to be *individually recoverable* if this change in conditional distribution is not obtainable from other combinations of features that arise with reasonable probability. This statistical property is hard to work with and below we suggest some (possibly too strong) combinatorial properties of the support of A that imply it. Better formalizations seem quite plausible and are left for future work.

2 Definitions and Results

The dictionary is an unknown matrix $A \in \mathbb{R}^{n \times m}$. We are given i.i.d samples y that are generated by $y = Ax$, where $x \in \mathbb{R}^m$ is chosen from some distribution. We have N samples $y^i = Ax^i$ for $i = 1, \dots, N$. As in the introduction, we will refer to coordinates of x as *features* and those of y as *pixels*, even though vision isn’t the only intended application. For most of the paper we assume the entries of x are independent Bernoulli variables: each x_i is 1 with probability ρ and 0 with probability $1 - \rho$; we refer to this as ρ -Bernoulli distribution. This assumption can be relaxed somewhat: we only need that entries of x are pairwise independent, and $e^T x$ should satisfy concentration bounds for reasonable vectors e . The nonzero entries of x can also be in $[1, c]$ instead of being exactly 1.

The j th column of A is denoted by A_j , the i th row of A by $A^{(i)}$, and the entries of A are denoted by $A_j^{(i)}$.

For ease of exposition we first describe our learning algorithm for nonnegative dictionaries (i.e., $A_j^{(i)} \geq 0$, for all $i \in [n], j \in [m]$) and then in Section 4 describe the generalization to the general case. Note that the subcase of nonnegative dictionaries is also of practical interest.

2.1 Nonnegative dictionaries

By normalizing, we can assume without loss of generality that the expected value of each pixel is 1, that is, $\mathbb{E}[y_i] = \mathbb{E}[(Ax)_i] = 1$. We also assume that $|A_j^{(i)}| \leq \Lambda$ for some constant Λ

²: no entry of A is too large. Let G_b be the bipartite graph defined by entries of A that have magnitudes larger than or equal to b , that is, $G_b = \{(i, j) : |A_j^{(i)}| \geq b, i \in [n], j \in [m]\}$. We make two assumptions about this graph (the parameters d, σ and κ will be chosen later).

Assumption 1: (Every feature has significant effect on pixels) There are at least d edges with weights larger than σ for every feature j . That is, the degree of G_σ on the feature side is always larger than d .

Assumption 2: (Low pairwise intersections among features) In G_τ the neighborhood of each feature (that is, $\{i \in [n] : A_j^{(i)} \geq \tau\}$) has intersection up to $d/10$ (with total weight $< d\sigma/10$) with each of at most $o(1/\sqrt{\rho})$ other features, and intersection at most κ with the neighborhood of each remaining features. Here τ is $O_\theta(1/\log n)$ as explained below and $\kappa = O_\theta(d/\log^2 n)$.

Guideline through notation: We will think of $\sigma \leq 1$ as a small constant, $\Lambda \geq 1$ a constant, and Δ a sufficiently large constant which is used to control the assumption. Let $\theta = (\sigma, \Lambda, \Delta)$ and we use the notation $O_\theta(\cdot)$ to hide the dependencies of σ, Λ, Δ . Also, we think of m as not much larger than n , $\rho < 1/\text{poly}(\log n)$, and $d \ll n$. The normalization assumption implies (for all practical purposes in the algorithm) that $md\rho \in [n/\Lambda, n/\tau]$. We typically think of d as $1/\rho$, hence a running time of m^d would be bad (though it is unclear a priori how to even achieve that).

Precisely, for our algorithms to work, we need $d \geq \Delta\Lambda \log^2 n/\sigma^2$, $\tau = O(\sigma^4/\Delta\Lambda^2 \log n) = O_\theta(1/\log n)$ and $\kappa = O(\sigma^8 d/\log^2 n \Delta^2 \Lambda^6) = O_\theta(d/\log^2 n)$ for some sufficiently large constant Δ and the density $\rho = o(\sigma^5/\Lambda^{6.5} \log^{2.5} n) = o_\theta(1/\log^{2.5} n)$. Note that if G_τ were like a random graph (i.e. if features affect random sets of d pixels) when $d^2 \ll n$, the pairwise intersection κ between the neighborhoods of two features in G_τ would be $O(1)$. However, we allow these intersections to be $\kappa = O_\theta(d/\log^2 n)$.

Now we give a stronger version of Assumption 2 which will allow a stronger algorithm.

Assumption 2': In G_τ , the pairwise intersection of the neighborhoods of any two features j, k is less than κ , where $\tau = O_\theta(1/\log n)$ and $\kappa = O_\theta(d/\log^2 n)$.

The algorithm can only learn the real-valued matrix approximately. Two dictionaries are *close* if they satisfy the following definition:

DEFINITION 1 (ϵ -EQUIVALENT) *Two dictionaries A and $\hat{A} \in R^{n \times m}$ are ϵ -equivalent, if for a random vector $x \in \mathbb{R}^m$ with independent ρ -Bernoulli components, with high probability Ax and $\hat{A}x$ are entry-wise ϵ -close.*

THEOREM 1 (NONNEG CASE)

Under Assumptions 1 and 2, when $\rho = o(\sigma^5/\Lambda^{6.5} \log^{2.5} n) = o_\theta(1/\log^{2.5} n)$, Algorithm 2 runs in $n^{O(\Lambda \log^2 n/\sigma^4)}$ time, uses $\text{poly}(n)$ samples and outputs a matrix that is $o(\rho)$ -equivalent to the true dictionary A . Furthermore, under Assumptions 1 and 2' the same algorithm returns a dictionary that is n^{-C} -equivalent to the true dictionary, while using n^{4C+3} samples, where C is a large constant depending on Δ .³

²Though the absolute value notations are redundant in the nonnegative case, we keep them so that they can be adapted to the general case. Similarly for the definition of G_b later.

³Recall that Δ is a sufficiently large constant that controls the parameters of the assumptions.

The theorem is proved in Section 3.

Remark: Assumption 2 is *existentially* close to optimal, in the sense that if it is significantly violated: e.g., if there are $\text{poly}(1/\rho)$ features that intersect the neighborhood of feature j using edges of total weight $\Omega(\ell_1\text{-norm of } A_j)$ then feature j is no longer individually recoverable: its effect can be duplicated whp by combinations of these other features. But a more precise characterization of individual recoverability would be nice, as well as a matching algorithm.

2.2 Dictionaries with negative entries

When the edges can be both positive and negative, it is no longer valid to assume the expectation of y_i 's are equal to 1. Instead, we choose a different normalization: the variances of y_i 's are 1. We still assume magnitude of edge weights are at most Λ , and that features don't overlap a lot as described in Assumption 1 and 2. We also need one more assumption to bound the variance contributed by the small entries.

Assumption G1: The degree of G_σ on side of x is always larger than $2d$.

Assumption G2': In G_τ , the pairwise intersection of the neighborhoods of any two features j, k is less than κ , where $\tau = O_\theta(1/\log n)$ and $\kappa = O_\theta(d/\log^2 n)$.

Assumption G3: (small entries of A don't cause large effects) $\rho \|A_{\leq \tau}^{(i)}\|_2^2 \leq \gamma$, where $A_{\leq \delta}^{(i)}$ be the vector that only contains the entries of $A^{(i)}$ that are at most δ , and $\gamma = \sigma^4/2\Delta\Lambda^2 \log n$.

Note that Assumption G1 differs from Assumption 1 by a constant factor 2 just to simplify some notations later. Assumption G2' is the same as before.

This assumption G3 intuitively says that for each $y_i = \sum_k A_k^{(i)} x_k$, the smaller $A_k^{(i)}$'s should not contribute too much to the variance of y_i . This is automatically satisfied for nonnegative dictionaries because there can be no cancellations. Notice that this assumption is talking about rows of matrix A (corresponding to pixels), whereas the earlier assumptions talk about columns of A (corresponding to features). Also, consider τ to be the smallest number between what is required by Assumption G2' and G3.

In term of parameters, we still need $d \geq \Delta\Lambda \log^2 n/\sigma^2$, and $\kappa = O(\sigma^8 d/\log^2 n \Delta^2 \Lambda^6) = O(d/\log^2 n)$. As before Δ is a large enough constant.

THEOREM 2

Under Assumptions G1, G2' and G3, when $\rho = o(\sigma^5/\Lambda^{6.5} \log^{2.5} n) = o_\theta(1/\log^{2.5} n)$ there is an algorithm that runs in $n^{O(\Delta\Lambda \log^2 n/\sigma^2)}$ time, uses $n^{4C+5}m$ samples and outputs a matrix that is n^{-C} -equivalent to the true dictionary A , where C is a constant depending on Δ .

The algorithm and the proof of Theorem 2 are sketched in Section 4.

3 Nonnegative Dictionary Learning

Recall that dictionary learning seems hard because both A and x are unknown. To get around this problem, previous works (e.g. [AGM13]) try to extract information about the

assignment x without first learning A (but assuming nice properties of A). After finding x , recovering A becomes easy. In [AGM13] the unknown x 's were recovered via an overlapping clustering procedure. The procedure relies on incoherence of A , as when A is incoherent it is possible to test whether the support of x^1, x^2 intersect. This idea fails when x is only slightly sparse, because in this setting the supports of x^1, x^2 always have a large intersection.

Our algorithm here relies on correlation among *pixels*. The key observation is: if the j th bit in x is 1, then $Ax = A_j + \sum_{k \neq j} A_k x_k$. Pixels with high values in A_j tend to be *elevated* above their mean values (recall A is nonnegative). At first it is unclear how this simultaneous elevation can be spotted, since A_j is unknown and these elevations/correlations among pixels are much smaller than the standard deviation of individual pixels. Therefore we look for local regions —small subsets of pixels— in A_j where this effect is significant in the aggregate (i.e., sum of pixel values), and can be used to consistently predict the value of x_j . These are called the *signature sets* (see Definition 2). If we can identify signature sets, they can give us a good estimation of whether the feature x_j is present.

Since the signature sets are small, in quasi-polynomial time we can afford to enumerate all sets of that size, and check if the pixels in these sets are likely to be elevated together. However, this does not solve the problem, because there can be many sets —called *correlated sets* below— that show similar *correlations* and look similar to signature sets. It is hard to separate signature sets from other correlated sets when the size of the set is small. This leads to the next idea: try to *expand* a signature set by first estimating the corresponding column of A , and then picking large entries in that column. The resulting sets are called *expanded signature sets*; these have size d (and hence could not have been found by exhaustive guessing alone). If the set being expanded is indeed a signature set, this expansion process can correctly estimate the column of A . We give algorithms that can find expanded signature sets, and using these sets we can get a rough estimation for the matrix A . Finally, we also give a procedure that leverages the *individually recoverable* properties of the features, and refines the solution to be inverse polynomially equivalent to the true dictionary.

The high level algorithm is described in Algorithm 2 (the concepts such as correlated sets, and empirical bias are defined later). To simplify the proof the algorithm description uses Assumption 2'; we summarize later (in Section 3.4) what changes with Assumption 2.

The main algorithm has three main steps. Section 3.1 explains how to test for correlated sets and expand a set (1-2 in Algorithm 2); Section 3.2 shows how to find expanded signature sets and a rough estimation of A (3-6 in Algorithm 2); finally Section 3.3 shows how to refine the solution and get \hat{A} that is inverse polynomially equivalent to A (7-10 in Algorithm 2).

3.1 Correlated Sets, Signature Sets and Expanded Sets

We consider a set T of size $t = \Omega(\text{poly log } n)$ (to be specified later), and denote by β_T the random variable representing the sum of all pixels in T , i.e., $\beta_T = \sum_{i \in T} y_i$. We can expand β_T as

$$\beta_T = \sum_{i \in T} y_i = \sum_{i \in T} \left(\sum_{j=1}^m A_j^{(i)} x_j \right) = \sum_{j=1}^m \left(\sum_{i \in T} A_j^{(i)} \right) x_j.$$

Let $\beta_{j,T} = \left(\sum_{i \in T} A_j^{(i)} \right)$ be the contribution of x_j to the sum β_T , then β_T is just

$$\beta_T = \sum_{j=1}^m \beta_{j,T} x_j \quad (1)$$

Note that by the normalization of $\mathbb{E}[y_i]$, we have $\mathbb{E}[\beta_T] = \sum_{i \in T} \mathbb{E}[y_i] = t$. Intuitively, if for all j , $\beta_{j,T}$'s are relatively small, β_T should concentrate around its mean. On the other hand, if there is some j whose coefficient $\beta_{j,T}$ is significantly larger than other $\beta_{k,T}$, then β_T will be *elevated* by $\beta_{j,T}$ precisely when $x_j = 1$. That is, with probability roughly ρ (corresponding to when $x_j = 1$), we should observe β_T to be roughly $\beta_{j,T}$ larger than its expectation.

Now we make this precise by defining such set T with only one large coefficient $\beta_{k,T}$ as signature sets.

DEFINITION 2 (SIGNATURE SET) *A set T of size t is a signature set for x_j , if $\beta_{j,T} \geq \sigma t$, and for all $k \neq j$, the contribution $\beta_{k,T} \leq \sigma^2 t / \Delta \log n$. Here Δ is a large enough constant.*

The following lemma formalizes the earlier intuition that if T is a signature set for x_j , then a large β_T is highly correlated with the event $x_j = 1$.

LEMMA 3

Suppose T of size t is a signature set for x_j with $t = \omega(\sqrt{\log n})$. Let E_1 be the event that $x_j = 1$ and E_2 be the event that $\beta_T \geq \mathbb{E}[\beta_T] + 0.9\sigma t$. Then for large constant C (depending on the Δ in Definition 2)

1. $\Pr[E_1] + n^{-2C} \geq \Pr[E_2] \geq \Pr[E_1] - n^{-2C}$.
2. $\Pr[E_2|E_1] \geq 1 - n^{-2C}$, and $\Pr[E_2|E_1^c] \leq n^{-2C}$.
3. $\Pr[E_1|E_2] \geq 1 - n^{-C}$.

PROOF: We can write β_T as

$$\beta_T = \beta_{j,T} x_j + \sum_{k \neq j} \beta_{k,T} x_k \quad (2)$$

The idea is that since $\beta_{k,T} < \sigma^2 t / (\Delta \log n)$ for all $k \neq j$, the summation in the RHS above is highly concentrated around its mean, which is actually very close to $\mathbb{E}[\beta_T] = t$. Therefore since $\beta_{j,T} > \sigma t$, we know $\beta_T > t + 0.9\sigma t$ essentially iff $x_j = 1$.

Formally, observe that $\mathbb{E}[\beta_{j,T} x_j] = \rho \beta_{j,T} \leq \rho \Delta t = o(\sigma t)$, and recall that $\mathbb{E}[\beta_T] = t$, we have $\mathbb{E}[\sum_{k \neq j} \beta_{k,T} x_k] = (1 - o(\sigma))t$. Let $M = \sigma^2 t / (\Delta \log n)$ be the upper bound for $\beta_{k,T}$, and then the variance of the sum $\sum_{k \neq j} \beta_{k,T} x_k$ is bounded by $\rho M \sum_{k \neq j} \beta_{k,T} \leq M t$. Then by calling Bernstein inequality (see Theorem 23, but note that σ there is the standard deviation), we have

$$\Pr \left[\left| \sum_{k \neq j} \beta_{k,T} x_k - \mathbb{E}[\sum_{k \neq j} \beta_{k,T} x_k] \right| > \sigma t / 20 \right] \leq 2 \exp \left(- \frac{\sigma^2 t^2 / 400}{2Mt + \frac{2}{3} \frac{M}{\sqrt{Mt}} \sigma t / 20} \right) \leq n^{-2C}.$$

where C is a large constant depending Δ .

Part (2) immediately follows: if $x_j = 1$, then $\beta_T < t + 0.9\sigma t$ iff the sum deviates from its expectation by more than $\sigma t/20$, which happens with probability $< n^{-2C}$. So also if $x_j = 0$, E_2 occurs with probability $< n^{-2C}$.

This then implies part (1), since the probability of E_1 is precisely ρ .

Combining the (1) and (2), and using Bayes' rule $\Pr[E_1|E_2] = \Pr[E_2|E_1] \Pr[E_1] / \Pr[E_2]$, we obtain (3). \square

Thus if we can find a signature set for x_j , we would roughly know the samples in which $x_j = 1$. The following lemma shows that assuming the low pairwise assumptions among features, there *exists* a signature set for every feature x_j .

LEMMA 4

Suppose A satisfies Assumptions 1 and 2, let $t = \Omega(\Lambda \Delta \log^2 n / \sigma^2)$, then for any $j \in [n]$, there exists a signature set of size t for node x_j .

PROOF: We show the existence by probabilistic method. By Assumption 1, node x_j has at least d neighbors in G_σ . Let T be a uniformly random set of t neighbors of x_j in G_σ . Now by the definition of G_σ we have $\beta_{j,T} \geq \sigma t$.

Using a bound on intersection size (Assumption 2') followed by Chernoff bound, we show that T is a signature set with good probability. For $k \neq j$, let $f_{k,T}$ be the number of edges from x_k to T in graph G_τ . Then we can upperbound $\beta_{k,T}$ by $t\tau + f_{k,T}\Lambda$ since all edge weights are at most Λ and there are at most $f_{j,T}$ edges with weights larger than τ . Using simple Chernoff bound and union bound, we know that with probability at least $1 - 1/n$, for all $k \neq j$, $f_{k,T} \leq 4 \log n$. Therefore $\beta_{k,T} \leq t\tau + f_{k,T}\Lambda \leq \sigma^2 t / (\Delta \log n)$ for $t \geq \Omega(\Lambda \Delta \log^2 n / \sigma^2)$, and $\tau = O(\sigma^2 / \Delta \log n)$. \square

Although signature sets exist for all x_j , it is difficult to find them; even if we enumerate all subsets of size t , it is not clear how to know when we found a signature set. Thus we first look for ‘‘correlated’’ sets, which are defined as follows:

DEFINITION 3 (CORRELATED SET) *A set T of size t is called correlated, if with probability at least $\rho - 1/n^2$ over the choice of x 's, $\beta_T \geq \mathbb{E}[\beta_T] + 0.9\sigma t = t + 0.9\sigma t$.*

It follows easily (Lemma 3) that signature sets must be correlated sets.

COROLLARY 5

If T of size t is a signature set for x_j , and $t = \omega(\sqrt{\log n})$, then T is a correlated set.

Although signature sets are all correlated sets, the other direction is far from true. There can be many correlated sets that are not signature sets. A simple counterexample would be that there are j and j' such that both $\beta_{j,T}$ and $\beta_{j',T}$ are larger than σt . This kind of counterexample seems inevitable for any test on a set T of polylogarithmic size.

To resolve this issue, the idea is to *expand* any set of size t into a much larger set \tilde{T} , which is called *expanded set for T* . If T happened to be a signature set to start with, such an expansion would give good estimate of the corresponding column of A , and more importantly, \tilde{T} will have a similar ‘signature’ property as T , which we can now verify because \tilde{T} is large.

Algorithm 1 and Definition 4 show how to expand T to \tilde{T} . The empirical expectation $\hat{\mathbb{E}}[f(y)]$ is defined to be $\frac{1}{N} \sum_{i=1}^N f(y^i)$.

Algorithm 1. $\tilde{T} = \text{expand}(T, \text{threshold})$

Input: T of size t , d , and N samples $\{y^1, \dots, y^N\}$

Output: vector $\tilde{A}_T \in \mathbb{R}^n$ and expanded set \tilde{T} of size d (when T is a signature set \tilde{A}_T is an estimation of A_j).

- 1: *Recovery Step:* Let L be the set of samples whose β_T values are larger than $\hat{\mathbb{E}}[\beta_T] + \text{threshold}$

$$L = \left\{ y^k \mid \beta_T^k \geq \hat{\mathbb{E}}[\beta_T] + \text{threshold} \right\}$$

- 2: *Estimation Step:* Compute the empirical mean of samples in L , and obtain \tilde{A}_T by shifting and scaling

$$\hat{\mathbb{E}}_L[y] = \frac{1}{|L|} \left(\sum_{y^k \in L} y^k \right), \text{ and } \tilde{A}_T(i) = \max\{0, (\hat{\mathbb{E}}_L[y_i] - \hat{\mathbb{E}}[y_i]) / (1 - \rho)\}$$

- 3: *Expansion Step:* $\tilde{T} = \{d \text{ largest coordinates of } \tilde{A}_T\}$
-

DEFINITION 4 For any set T of size t , the expanded set \tilde{T} for T is defined as the one output by Algorithm 1. The estimation \tilde{A}_T is the output at step 2.

When T is a signature set for x_j , then \tilde{A}_T is already close to the true A_j , and the expanded set \tilde{T} is close to the largest entries of A_j .

LEMMA 6

If T is a signature set for x_j and the number of samples $N = \Omega(n^{2+\delta}/\rho^3)$, where δ is any positive constant, then with high probability $\|\tilde{A}_T - A_j\|_\infty \leq 1/n$. Furthermore, $\beta_{j, \tilde{T}}$ is d/n -close to the sum of d largest entries in A_j , and for all $i \in \tilde{T}$, $A_j^{(i)} \geq 0.9\sigma$.

PROOF: Let's first consider $\mathbb{E}[\tilde{A}_T] \triangleq (\mathbb{E}[y|E_2] - \mathbf{1}) / (1 - \rho)$ where E_2 is the event that $\beta_T \geq t + 0.9\sigma t$ defined in Lemma 3. Recall that because of normalization, we know for any j , $\sum_{i \in [n]} A_j^{(i)} = 1/\rho$, so in particular $y_i \leq 1/\rho$. By Lemma 3 and some calculations (see Lemma 21), we have that $\|\mathbb{E}[y|E_2] - \mathbb{E}[y|E_1]\|_\infty \leq n^{-C}/\rho$. Note that $\mathbb{E}[y|E_1] = 1 + (1 - \rho)A_j$. Therefore we have that $\|\mathbb{E}[\tilde{A}_T] - A_j\|_\infty \leq n^{-C}/\rho$.

Now by concentration inequalities when $N = \Omega(n^{2+\delta}/\rho^3)$ (notice that the variance of each coordinate is bounded by Λ), $\|\tilde{A}_T - \mathbb{E}[\tilde{A}_T]\|_\infty \leq 1/n$ with very high probability ($\exp(-\Omega(n^\delta))$). This probability is high enough so we can apply union bound for all signature sets. \square

3.2 Identify Expanded Signature Sets

We will now see the advantage that the expanded sets \tilde{T} provide. If T happens to be a signature set, the expanded set \tilde{T} for T also has similar property. But now \tilde{T} is a much larger set (size d as opposed to $t = \text{polylog}$), and we know (by Assumption 2') that different features have limited intersection, so if we see a large elevation it is likely to be caused by

a single feature! We will leverage this in order to identify expanded signature sets among all the expanded sets.

If an expanded set \tilde{T} also has essentially a unique large coefficient $\beta_{j,\tilde{T}}$, we call it an *expanded signature set*.

DEFINITION 5 (EXPANDED SIGNATURE SET) *An expanded set \tilde{T} is an expanded signature set for x_j if $\beta_{j,\tilde{T}} \geq 0.7\sigma d$ and for all $k \neq j$, $\beta_{k,\tilde{T}} \leq 0.3\sigma d$.*

Note that an expanded signature set always has size d and the gap between largest $\beta_{j,T}$ and the second largest is only constant as opposed to logarithmic in the definition of signature set. As its name suggests, an expanded set \tilde{T} of a signature set T for x_j is an expanded signature set for x_j as well. On one hand, the Lemma 6 guarantees that \tilde{T} connects to x_j with large weights, and on the other hand, since the pairwise intersection of neighborhoods of x_j and x_k in G_τ is small, \tilde{T} cannot also connect to other x_k with too many large weights.

LEMMA 7

If T is a signature set for x_j , then the expanded set \tilde{T} for T is always an expanded signature set for x_j . In fact, the coefficient $\beta_{j,\tilde{T}}$ is at least $0.9\sigma d$.

PROOF: Since we know there are at least d weights $A_j^{(i)}$ bigger than σ for any column A_j , by Lemma 6 we know $\beta_{j,\tilde{T}} \geq \sigma d - o(1)d \geq 0.9\sigma d$.

Furthermore, Lemma 6 says x_j connects to every node in \tilde{T} with weights larger than 0.9σ (since by Assumption 1 there are more than d edges of weight at least σ from node j). By Assumption 2 on the graph, for any other $k \neq j$, the number of y_i 's that are connected to both k and j in G_τ is bounded by κ . In particular, the number of edges from k to \tilde{T} with weights more than τ is bounded by κ . Therefore the coefficient $\beta_{k,\tilde{T}} = \sum_{(i,k) \in G_\tau} A_k^{(i)} + \sum_{(i,k) \notin G_\tau} A_k^{(i)}$ is bounded by $\Lambda\kappa + |\tilde{T}|\tau = o(d) \leq 0.3d$. (Recall $\tau = o(1)$ and $\kappa = o(d)$) \square

The following notion of empirical bias is a more precise way (compared to correlated set) to measure the simultaneous elevation effect.

DEFINITION 6 (EMPIRICAL BIAS) *The empirical bias $\hat{B}_{\tilde{T}}$ of an expanded set \tilde{T} of size d is defined to be the largest B that satisfies*

$$\left| \left\{ k \in [p] : \beta_{\tilde{T}}^k \geq \hat{\mathbb{E}}[\beta_{\tilde{T}}] + B \right\} \right| \geq \rho N/2.$$

In other words, $\hat{B}_{\tilde{T}}$ is the difference between the $\rho N/2$ -th largest $\beta_{\tilde{T}}^k$ in the samples and $\hat{\mathbb{E}}[\beta_{\tilde{T}}]$.

The key lemma in this part shows the expanded set with largest empirical bias must be an expanded signature set:

LEMMA 8

Let \tilde{T}^ be the set with largest empirical bias $\hat{B}_{\tilde{T}^*}$ among all the expanded sets \tilde{T} . The set \tilde{T}^* is an expanded signature set for some x_j .*

We build this lemma in several steps. First of all, we show that the bias of \tilde{T} is almost equal to the largest $\beta_{j,\tilde{T}}$: if $\beta_{\tilde{T}}$ contains a large term $\beta_{j,\tilde{T}}x_j$, then certainly this term will contribute to the bias $\hat{B}_{\tilde{T}}$; on the other hand, suppose in some extreme case $\beta_{\tilde{T}}$ only has two non-zero terms $\beta_{j,\tilde{T}}x_j + \beta_{k,\tilde{T}}x_k$. Then they cannot contribute more than $\max\{\beta_{j,\tilde{T}}, \beta_{k,\tilde{T}}\}$ to the bias, because otherwise both x_k and x_j have to be 1 to make the sum larger than $\max\{\beta_{j,\tilde{T}}, \beta_{k,\tilde{T}}\}$, and this only happens with small probability $\rho^2 \ll \rho$.

The intuitive argument above is not far from true: basically we could show that a) There are indeed very few large coefficients $\beta_{k,\tilde{T}}$'s (see Claim 9 for the precise statement) b) the sum of those small $\beta_{k,\tilde{T}}x_k$ concentrates around its mean, thus won't contribute much to the bias.

After relating the bias of \tilde{T} to the largest coefficients $\max_j \beta_{j,\tilde{T}}$, we further argue that taking the set \tilde{T}^* with largest bias among all the \tilde{T} , we not only see a large coefficient $\beta_{j,\tilde{T}}$, but also we observe a gap between the top $\beta_{j,\tilde{T}}$ and all other $\beta_{k,\tilde{T}}$'s, and hence \tilde{T} is an expanded signature set for x_j .

We make the arguments above precise by the following claims. First, we shall show there cannot be too many large coefficients $\beta_{j,D}$ for any set D of size d (although we only apply the claim on expanded sets).

CLAIM 9

For any set \tilde{T} of size d , the number of k 's such that $\beta_{k,\tilde{T}}$'s is larger than $d\sigma^4/\Delta\Lambda^2 \log n$ is at most $O(\Delta\Lambda^3 \log n/\sigma^4)$.

PROOF: For the ease of exposition, we define $K_{large} = \{k : \beta_{k,\tilde{T}} \geq d\sigma^4/\Delta\Lambda^2 \log n\}$. Hence the goal is to prove that $|K_{large}| \leq O(\Delta\Lambda^3 \log n/\sigma^4)$. Recall that $\beta_{k,\tilde{T}} = \sum_{i \in \tilde{T}} A_k^{(i)}$. Let $Q_k = \{i \in \tilde{T} : A_k^{(i)} \geq \tau\}$ be the subset of nodes in \tilde{T} that connect to k with weights larger than τ . We have that $\beta_{k,\tilde{T}} = \sum_{i \notin Q_k} A_k^{(i)} + \sum_{i \in Q_k} A_k^{(i)}$. The first sum is upper bounded by $d\tau \leq d\sigma^4/2\Delta\Lambda^2 \log n$. Therefore for $k \in K_{large}$, the second sum is lower bounded by $d\sigma^4/2\Delta\Lambda^2 \log n$. Since $A_k^{(i)} \leq \Lambda$, we have $|Q_k| \geq \sigma^4 d/2\Delta\Lambda^3 \log n$.

On the other hand, by Assumption 2 we know in graph G_τ , any two features cannot share too many pixels: for any k and k' , $|Q_k \cap Q_{k'}| \leq \kappa$. Also note that by definition, $Q_j \subset \tilde{T}$, which implies that $|\cup_{k \in K_{large}} Q_k| \leq |\tilde{T}| = d$. By inclusion-exclusion we have

$$d \geq \left| \bigcup_{k \in K_{large}} Q_k \right| \geq \sum_{k \in K_{large}} |Q_k| - \sum_{k, k' \in K_{large}} |Q_k \cap Q_{k'}| \geq |K_{large}| \sigma^4 d/2\Delta\Lambda^3 \log n - |K_{large}|^2/2 \cdot \kappa \quad (3)$$

This implies that $|K_{large}| \leq O(\Delta\Lambda^3 \log n/\sigma^4)$, when $\kappa = O(\sigma^8 d/\Delta^2 \Lambda^6 \log^2 n)$. ⁴ \square

For simplicity, let $k^* = \arg \max_k \beta_{k,\tilde{T}}$, so $\beta_{k^*,\tilde{T}}$ is the largest coefficient in $\beta_{\tilde{T}}$. Recall that the definition of expanded signature set roughly translates to a constant factor gap between $\beta_{k^*,\tilde{T}}$ to any other coefficient $\beta_{k,\tilde{T}}$.

The next claim shows that the empirical bias $\hat{B}_{\tilde{T}}$ is a good estimate of $\beta_{k^*,\tilde{T}}$ when $\beta_{k^*,\tilde{T}}$ is large.

⁴Note that any subset of K_{large} also satisfies equation (3), thus we don't have to worry about the other range of the solution of (3)

CLAIM 10

For any expanded \tilde{T} of size d , with high probability over the choices of all the N samples, the empirical bias $\hat{B}_{\tilde{T}}$ is within $0.1d\sigma^2/\Lambda$ to $\beta_{k^*,\tilde{T}} = \max_k \beta_{k,\tilde{T}}$ when $\beta_{k^*,\tilde{T}}$ is at least $0.5d\sigma$.

PROOF: Let $K'_{large} = K_{large} \setminus \{k^*\}$ ⁵, and $\beta_{small,\tilde{T}} = \sum_{k \notin K_{large}} \beta_{k,\tilde{T}} x_k$, and $\beta_{large,\tilde{T}} = \sum_{k \in K'_{large}} \beta_{k,\tilde{T}} x_k$.

First of all, the variance of $\beta_{small,\tilde{T}}$ is bounded by $\rho \sum_{k \notin K_{large}} \beta_{k,\tilde{T}}^2 \leq d\sigma^4/\Delta\Lambda^2 \log n$. $\left(\rho \sum_{k \notin K_{large}} \beta_{k,\tilde{T}}\right) \leq d^2\sigma^4/\Delta\Lambda^2 \log n$. By Bernstein's inequality, for sufficiently large Δ , with probability at most $1/n^2$ over the choice of x , the value $|\beta_{small,\tilde{T}} - \mathbb{E}[\beta_{small,\tilde{T}}]|$ is larger than $0.05d\sigma^2/\Lambda$, that is, $\beta_{small,\tilde{T}}$ nicely concentrates around its mean. Secondly, with probability at most ρ we have $x_{k^*} = 1$, and then $\beta_{k^*,\tilde{T}} x_{k^*}$ is elevated above its mean by roughly $\beta_{k^*,\tilde{T}}$. Thirdly, the mean of $\beta_{large,\tilde{T}}$ is at most $\rho \sum_{k \in K'_{large}} \beta_{k,\tilde{T}} \leq \rho|K|d$, which is $o(\sigma d)$ by Claim 9. These three points altogether imply that with probability at least $\rho - n^{-2}$, $\beta_{\tilde{T}}$ is above its mean by $\beta_{k^*,\tilde{T}} - 0.1\sigma^2 d/\Lambda$. Also note that the empirical mean $\hat{\mathbb{E}}[\beta_{\tilde{T}}]$ is sufficiently close to the $\beta_{\tilde{T}}$ with probability $1 - \exp(-\Omega(n))$ over the choices of N samples, when $N = \text{poly}(n)$. Therefore with probability $1 - \exp(-\Omega(n))$ over the choices of N samples, $\hat{B}_{\tilde{T}} > \beta_{k^*,\tilde{T}} - 0.1\sigma^2 d/\Lambda$.

It remains to prove the other side of the inequality, that is, $\hat{B}_{\tilde{T}} \leq \beta_{k^*,\tilde{T}} + 0.1\sigma^2 d/\Lambda$.

Note that $|K_{large}| = O(\log n)$, thus with probability at least $1 - 2\rho^2|K|^2$, at most one of the x_k , ($k \in K_{large}$) is equal to 1. Then with probability at least $1 - 2\rho^2|K|^2$ over the choices of x , $\beta_{large,\tilde{T}} + \beta_{k^*,\tilde{T}}$ is elevated above its mean by at most $\beta_{k^*,\tilde{T}}$. Also with probability $1 - n^{-2}$ over the choices of x , $\beta_{small,\tilde{T}}$ is above its mean by at most $0.1\sigma^2 d/\Lambda$. Therefore with probability at least $1 - 3\rho^2|K|^2$ over the choices of x , $\beta_{\tilde{T}}$ is above its mean by at most $\beta_{k^*,\tilde{T}} + 0.1\sigma d/\Lambda$. Hence when $3\rho^2|K|^2 \leq \rho/3$, with probability at least $1 - \exp(-\Omega(n))$ over the choice of the N samples, $\hat{B}_{\tilde{T}} \leq \beta_{k^*,\tilde{T}} + 0.1\sigma^2 d/\Lambda$. The condition is satisfied when $\rho \leq c/\log^2 n$ for a small enough constant c . \square

Now we are ready to prove Lemma 8.

PROOF:[of Lemma 8] By Claim 10 and the existence of good expanded signature sets (Lemma 7), we know the maximum bias is at least $0.8\sigma d$. Apply Claim 10 again, we know for the set \tilde{T}^* that has largest bias, there must be a feature j with $\beta_{j,\tilde{T}^*} \geq 0.7\sigma d$.

For the sake of contradiction, now we assume that the set \tilde{T}^* with largest bias is not an expanded signature set. Then there must be some $k \neq j$ where $\beta_{k,\tilde{T}^*} \geq 0.3\sigma d$. Let Q_j and Q_k be the set of nodes in \tilde{T}^* that are connected to j and k in G_τ (these are the same Q 's as in the proof of Claim 9). We know $|Q_j \cap Q_k| \leq \kappa$ by assumption, and $|Q_k| \geq 0.3\sigma d/\Lambda$. This means $|Q_j| \leq d - 0.3\sigma d/\Lambda + \kappa$ by inclusion-exclusion.

Now let T' be a signature set for x_j , and let \tilde{T}' be its expanded set, from Lemma 6 we know $\beta_{j,\tilde{T}'}$ is almost equal to the sum of the d largest entries in A_j , which is at least $0.2\sigma^2 d/\Lambda$ larger than β_{j,\tilde{T}^*} , since $|Q_j| \leq d - 0.2\sigma d/\Lambda$. By Claim 10 we know $\hat{B}(\tilde{T}') \geq \beta_{j,\tilde{T}'} - 0.1\sigma^2 d/\Lambda > \beta_{j,\tilde{T}^*} + 0.1\sigma^2 d/\Lambda \geq \hat{B}(\tilde{T}^*)$, which contradict with the assumption that \tilde{T}^* is the set with largest bias. \square

⁵ K_{large} is defined in proof of Claim 9

Now we have found expanded signature sets, we can then apply Algorithm 1 (but with threshold $0.6\sigma d$ instead of $0.9\sigma d$) on that to get an estimation.

LEMMA 11

If \tilde{T} is an expanded signature set for x_j , and $\tilde{A}_{\tilde{T}}$ is the corresponding column output by Algorithm 1, then with high probability $\|\tilde{A}_{\tilde{T}} - A_j\|_\infty \leq O(\rho(\Lambda^3 \log n / \sigma^2)^2 \sqrt{\Lambda \log n}) = o(\sigma)$.

PROOF: Define E_1 to be the event that $x_j = 1$, and E_2 to be the event that $\beta_{\tilde{T}} \geq 0.6d\sigma$.

When E_1 happens, event E_2 always happen unless $\beta_{\tilde{T}, \text{small}}$ is far from its expectation. In the proof of Claim 10 we've already shown the number of such samples is at most n with very high probability.

Suppose E_2 happens, and E_1 does not happen. Then either $\beta_{\tilde{T}, \text{small}}$ is far from its expectation, or at least two x_j 's with large coefficients $\beta_{j, \tilde{T}}$'s are on. Recall by Claim 9 the number of x_j 's with large coefficients is $|K| \leq O(\Lambda^3 \log n / \sigma^2)$, so the probability that at least two large coefficient is "on" (with $x_j = 1$) is bounded by $O(\rho^2 \cdot |K|^2) = \rho \cdot O(\rho \Lambda^6 \log^2 n / \sigma^4) = \rho \cdot o(\sigma / \sqrt{\Lambda \log n})$. With very high probability the number of such samples is bounded by $\rho N \cdot o(\sigma / \sqrt{\Lambda \log n})$.

Combining the two parts, we know the number of samples that is in $E_1 \oplus E_2$ (the symmetric difference between E_1 and E_2) is bounded by $\rho N \cdot o(\sigma / \sqrt{\Lambda \log n})$. Also, with high probability $(1 - n^{-C})$ all the samples have entries bounded by $O(\sqrt{\Lambda \log n})$ by Bernstein's inequality (variance of y_i is bounded by $\sum_j \rho(A_j^{(i)})^2 \leq \max_j A_j^{(i)} \sum_j \rho A_j^{(i)} \leq \Lambda$). Notice that this is a statement of the entire sample independent of the set T , so we do not need to apply union bound over all expanded signature sets.

Therefore by Lemma 21

$$\|\tilde{A}_{\tilde{T}} - A_j\|_\infty \leq o(\sigma / \sqrt{\Lambda \log n}) \cdot O(\sqrt{\Lambda \log n}) = o(\sigma).$$

□

The previous lemma looks very similar to the lemma for signature sets, however, the benefit is we know how to find a set that is guaranteed to be expanded signature set! So we can iteratively find all expanded signature sets.

After identifying $\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_k$ (reorder the columns of A to make them correspond to the first k columns), we can estimate the corresponding columns $\tilde{A}_{\tilde{T}_1}, \dots, \tilde{A}_{\tilde{T}_k}$. Since these are close to the true columns A_1, A_2, \dots, A_k (wlog. we reorder columns so $\tilde{A}_{\tilde{T}_j}$ correspond to A_j for $1 \leq j \leq k$), we can in fact compute $\hat{\beta}_{j, \tilde{T}} = \sum_{i \in \tilde{T}} \tilde{A}_{\tilde{T}_j}(i)$. By Lemma 11 we know $|\hat{\beta}_{j, \tilde{T}} - \beta_{j, \tilde{T}}| = o(\sigma d)$.

LEMMA 12

Having found \tilde{T}_i (and hence also $\tilde{A}_{\tilde{T}_i}$) for $i \leq k$, let \tilde{T} be the set with largest empirical bias among the expanded sets that have $\hat{\beta}_{j, \tilde{T}} < 0.2\sigma d$ for all $j \leq k$. Then \tilde{T} is an expanded signature set for new x_j where $j > k$.

PROOF: The proof is almost identical to Lemma 8.

First, if T is a signature set of x_j where $j > k$, then by Lemma 7 \tilde{T} must satisfy $\hat{\beta}_{j, \tilde{T}} < 0.2\sigma d$, so it will compete for the set with largest empirical bias.

Also, since $\hat{\beta}_{j, \tilde{T}} < 0.2\sigma d$, we know the coefficients in $\beta_{j, \tilde{T}}$ must have $j > k$. Leveraging this observation in the proof of Lemma 8 gives the result. □

3.3 Getting an Equivalent Dictionary

After finding expanded signature sets, we already have an estimation $\tilde{A}_{\tilde{T}_j}$ of A_j that is entry-wise $o(\sigma)$ close. However, this alone does not imply that the two dictionaries are ϵ -equivalent for very small ϵ .

In the final step, we look at *all* the large entries in the column A_j , and use them to identify whether feature x_j is 1 or 0. The ability to do this justifies the individually recoverable property of the dictionary.

LEMMA 13

Let S_j be the set of all entries larger than $\sigma/2$ in $\tilde{A}_{\tilde{T}_j}$, then $|S_j| \geq d$, $\beta_{j,S_j} \geq (0.5 - o(1)) |S_j| \sigma$, and for all $k \neq j$ $\beta_{k,S_j} \leq \sigma^2 |S_j| / \Delta \log n$ where Δ is a large enough constant.

PROOF: This follows directly from the assumptions. By Assumption 1, there are at least d entries in A_j that are larger than σ , all these entries will be at least $(1 - o(1))\sigma$ in $\tilde{A}_{\tilde{T}_j}$, so $|S_j| \geq d$.

Also, since for all $i \in S_j$, $\tilde{A}_{\tilde{T}_j}(i) \geq 0.5\sigma$, we know $A_j(i) \geq 0.5\sigma - o(\sigma)$, hence $\beta_{j,S_j} \geq (0.5 - o(1)) |S_j| \sigma$.

By Assumption 2, for any $k \neq j$, the number of edges in G_τ between k and S_j is bounded by κ , so $\beta_{k,S_j} \leq \tau |S_j| + \kappa \Lambda \leq \sigma^2 |S_j| / \Delta \log n$. \square

Since S_j has a unique large coefficient β_{j,S_j} , and the rest of the coefficients are much smaller, when Δ is large enough, and $N \geq n^{4C+\delta}/\rho^3$ we know \hat{A}_j is entry-wise $n^{-2C}/\log n$ close to A_j (this is using the same argument as in Lemma 6). We shall show this is enough to proof n^{-C} -equivalence between \hat{A} and A .

LEMMA 14

Let A, \hat{A} be dictionaries with rows having ℓ_1 -norm $O(1/\rho)$ and all entries in $A - \hat{A}$ have magnitude at most δ . Then \hat{A} and A are $O(\sqrt{\delta \log n})$ -equivalent.

The proof is an easy application of Bernstein's inequality (see Appendix A.1).

Remark: Notice that when $C \geq 1$ it is clear why \hat{A}_j should have ℓ_1 norm $1/\rho$ (because it is very close to A_j); when C is smaller we need to truncate the entries of \hat{A}_j that are smaller than $n^{-2C}/\log n$.

We now formally write down the steps in the algorithm.

3.4 Working with Assumption 2

In order to assume Assumption 2 instead of 2', we need to change the definition of signature sets to allow $o(1/\sqrt{\rho})$ "moderately large" ($\sigma t/10$) entries. This makes the definition look similar to expanded signature sets. Such signature sets still exist by similar probabilistic argument as in Lemma 4. Lemma 7 and Claims 9 and 10 can also be adapted.

Finally, for Lemma 14, the guarantee will be weaker (there can be $o(1/\sqrt{\rho})$ moderately large coefficients). The algorithm will only estimate x_j incorrectly if at least 6 such coefficients are "on" (has the corresponding x_j being 1), which happens with less than $o(\rho^3)$ probability. By argument similar to Lemma 6 and Lemma 14 we get the first part of Theorem 1.

Algorithm 2. Nonnegative Dictionary Learning

Input: N samples $\{y^1, \dots, y^N\}$ generated by $y^i = Ax^i$. Unknown dictionary A satisfies Assumptions 1 and 2.

Output: \hat{A} that is n^{-C} close to A

- 1: Enumerate all sets of size $t = O(\Lambda \log^2 n / \sigma^4)$, keep the sets that are correlated.
 - 2: Expand all correlated sets T , $\tilde{T} = \text{Expand}(T, 0.9\sigma t)$.
 - 3: **for** $j = 1$ **TO** m **do**
 - 4: Let \tilde{T}_j be the set with largest empirical bias, and for all $k < j$, $\hat{\beta}_{k, \tilde{T}} = \sum_{i \in \tilde{T}} \tilde{A}_{\tilde{T}_k}(i) \leq 2d\sigma$.
 - 5: Let $\tilde{A}_{\tilde{T}_k}$ be the result of estimation step in $\text{Expand}(\tilde{T}, 0.6\sigma d)$.
 - 6: **end for**
 - 7: **for** $j = 1$ **TO** m **do**
 - 8: Let S_j be the set of entries that are larger than $\sigma/2$ in $\tilde{A}_{\tilde{T}_j}$
 - 9: Let \hat{A}_i be the result of estimation step in $\text{Expand}(S_j, 0.4\sigma |S_j|)$
 - 10: **end for**
-

4 General Case

With minor modifications, our algorithm and its analysis can be adapted to the general case in which the matrix A can have both positive and negative entries.

We follow the outline from the non-negative case, and look at sets T of size t . The quantities β_T and $\beta_{j,T}$ are defined exactly the same as in Section 3.1. Additionally, let ν_T be the standard deviation of β_T , and let $\nu_{-j,T}$ be the standard deviation of $\beta_T - \beta_{j,T}x_j$. That is,

$$\nu_{-j,T}^2 = \mathbb{V}[\beta_T - \beta_{j,T}x_j] = \rho \sum_{k \neq j} \beta_{k,T}^2.$$

The definition of signature sets requires an additional condition to take into account the standard deviations.

DEFINITION 7 ((GENERAL) SIGNATURE SET) *A set T of size t is a signature set for x_j , if for some large constant Δ , we have: (a) $|\beta_{j,T}| \geq \sigma t$, (b) for all $k \neq j$, the contribution $|\beta_{k,T}| \leq \sigma^2 t / (\Delta \log n)$, and additionally, (c) $\nu_{-j,T} \leq \sigma t / \sqrt{\Delta \log n}$.*

In the nonnegative case the additional condition $\nu_{-j,T} \leq \sigma t / \sqrt{\Delta \log n}$ was automatically implied by nonnegativity and scaling. Now we use Assumption G3 to show there exist T in which (c) is true along with the other properties. To do that, we prove a simple lemma which lets us bound the variance (the same lemma is also used in other places).

LEMMA 15

Let T be a set of size t and S be an arbitrary subset of features, and consider the sum $\beta_{S,T} = \sum_{j \in S} \beta_{j,T}x_j$. Suppose for each $j \in S$, the number of edges from j to T in graph G_τ is bounded by W . Then the variance of $\beta_{S,T}$ is bounded by $2tW + 2t^2\gamma$.

PROOF: The idea is to split the weights $A_j^{(i)}$ into the *big* and *small* ones (threshold being τ). Intuitively, on one hand, the contribution to the variance from large weights is bounded

above because the number of such large edges is bounded by W . On the other hand, by assumption (3), the total variance of small weights is less than γ , which implies that the contribution of small weight to the variance is also bounded. Formally, we have

$$\begin{aligned}
\mathbb{V}[\beta_{S,T}] &= \rho \sum_{j \in S} \beta_{j,T}^2 = \rho \sum_{j \in S} \left(\sum_{i \in T} A_j^{(i)} \right)^2 \\
&= \rho \sum_{j \in S} \left(\sum_{i: i \in T, (i,j) \in G_\tau} A_j^{(i)} + \sum_{i: i \in T, (i,j) \notin G_\tau} A_j^{(i)} \right)^2 \\
&\leq 2\rho \sum_{j \in S} \left[\left(\sum_{i: i \in T, (i,j) \in G_\tau} A_j^{(i)} \right)^2 + \left(\sum_{i: i \in T, (i,j) \notin G_\tau} A_j^{(i)} \right)^2 \right] \\
&\leq 2\rho \sum_{j \in S} \left[W \left(\sum_{i: i \in T, (i,j) \in G_\tau} (A_j^{(i)})^2 \right) + t \left(\sum_{i: i \in T, (i,j) \notin G_\tau} (A_j^{(i)})^2 \right) \right] \\
&= 2\rho W \sum_{i \in T} \sum_{j \in S} (A_j^{(i)})^2 + 2\rho t \sum_{i \in T} \sum_{j: (i,j) \notin G_\tau} (A_j^{(i)})^2 \\
&\leq 2tW + 2t^2\gamma.
\end{aligned}$$

In the fourth line we used Cauchy-Schwarz inequality and in the last step, we used Assumption G3 about the total variance due to small terms being small, as well as the normalization of the variance in each pixel. \square

LEMMA 16

Suppose A satisfies our assumptions for general dictionaries, and let $t = \Omega(\Lambda \Delta \log^2 n / \sigma^2)$. Then for any $j \in [n]$, there exists a general signature set of size t for node x_j (as in Definition 7).

PROOF: As before, we use the probabilistic method. Suppose we fix some j . By Assumption G1, in G_σ , node x_j has either at least d positive neighbors or d negative ones. W.l.o.g., let us assume there are d negative neighbors. Let T be uniformly random subset of size t of these negative neighbors. By definition of G_σ , we have $\beta_{j,T} \leq -\sigma t$.

For $k \neq j$, let $f_{k,T}$ be the number of edges from x_k to T in graph G_τ . Using the same argument as in the proof of Lemma 4, we have $f_{k,T} \leq 4 \log n$ w.h.p. for all such $k \neq j$. Thus $|\beta_{k,T}| \leq t\tau + f_{k,T}\Lambda \leq \sigma^2 t / (\Delta \log n)$. Thus it remains to bound $\nu_{-j,T}$.

We could apply Lemma 15 with $W = 4 \log n \geq f_{k,T}$, and $S = [m] \setminus \{j\}$ on set T : we get $\nu_{-j,T}^2 \leq 2tW + 2t^2\gamma$. Recall that $\gamma = \sigma^2 / 3\Delta^2 \log n$ and thus $\nu_{-j,T} \leq \sigma t / \sqrt{\Delta \log n}$. \square

The proof of Lemma 3 now follows in the general case (here we will use the variance bound (c) in the general definition of signature sets), except that we need to redefine event E_2 to handle the negative case. For completeness, we state the general version of Lemma 3 in Appendix A.2. As before, signature sets give a great idea of whether $x_j = 1$.

Let us now define correlated sets: here we need to consider both positive and negative bias

DEFINITION 8 ((GENERAL) CORRELATED SET) *A set T of size t is correlated, if either with probability at least $\rho - 1/n^2$ over the choice of x 's, $\beta_T \geq \mathbb{E}[\beta_T] + 0.8\sigma t$, or with probability at least $\rho - 1/n^2$, $\beta_T \leq \mathbb{E}[\beta_T] - 0.8\sigma t$.*

Starting with a correlated set (a potential signature set), we expand it similar to (Definition 4), except that we find \tilde{T} as follows:

$$\tilde{T}_{temp} = \{2d \text{ coordinates of largest magnitude in } \hat{A}_T\}, \tilde{T}_1 = \{i \in \tilde{T}_{temp} : \hat{A}_T \geq 0\}$$

$$\tilde{T} = \begin{cases} \tilde{T}_1 & \text{if } |\tilde{T}_1| \geq d \\ \tilde{T}_{temp} \setminus \tilde{T}_1 & \text{otherwise} \end{cases}$$

Our earlier definitions of expanded signature sets and bias can also be adapted naturally:

DEFINITION 9 ((GENERAL) EXPANDED SIGNATURE SET) *An expanded set \tilde{T} is an expanded signature set for x_j if $|\beta_{j,\tilde{T}}| \geq 0.7\sigma d$ and for all $k \neq j$, $|\beta_{k,\tilde{T}}| \leq 0.3\sigma d$.*

Since Lemma 6 still holds, Lemma 7 follows straightforwardly. That is, there always exists a general expanded signature set \tilde{T} that is produced by a set T of size $t = O_\theta(\log n^2)$. (Note that this is why in the general case we assume that G_σ has degree at least $2d$ in Assumption G1. We want to make the size of good expanded set to be d instead of $d/2$ so that all the lemmas can be adapted without change of notation).

DEFINITION 10 ((GENERAL) EMPIRICAL BIAS) *The empirical bias $\hat{B}_{\tilde{T}}$ of an expanded set \tilde{T} of size d is defined to be the largest B that satisfies*

$$\left| \left\{ k \in [p] : \left| \beta_{\tilde{T}}^k - \hat{\mathbb{E}}[\beta_{\tilde{T}}] \right| \geq B \right\} \right| \geq \rho N/2.$$

In other words, $\hat{B}_{\tilde{T}}$ is the difference between the $\rho N/2$ -th largest $\beta_{\tilde{T}}^k$ in the samples and $\hat{\mathbb{E}}[\beta_{\tilde{T}}]$.

Let us now intuitively describe why the analog of Lemma 8 holds in the general case. We provide the formal statement and the proof in Appendix A.2

1. The first step, Claim 9 is a statement purely about the magnitudes of the edges (in fact, cancellations in $\beta_{k,\tilde{T}}$ for $k \neq j$ only help our case).
2. The second step, Claim 10 essentially argues that the small $\beta_{k,\tilde{T}}$ do not contribute much to the bias (a concentration bound, which still holds due to Lemma 15), and that the probability of *two* “large” features j, j' being on simultaneously is very small. The latter holds even if the $\beta_{j,\tilde{T}}$ have different signs.
3. The final step in the proof of Lemma 8 is an argument which uses the assumption on the overlap between features to contradict the maximality of bias, when the case where $\beta_{j,\tilde{T}}$ and $\beta_{j',\tilde{T}}$ are both “large”. This only uses the magnitudes of the entries in A , and thus also follows.

Recovering an equivalent dictionary. The main lemma in the nonnegative case, which shows that Algorithm 1 *roughly* recovers a column, is Lemma 11. The proof uses the property that signature sets are elevated “almost iff” the $x_j = 1$ to conclude that we get a good approximation to one of the columns. We have seen that this also holds in the general case, and since the rest of the argument deals only with the magnitudes of the entries, we conclude that we can roughly recover a column also in the general case. Let us state this formally.

LEMMA 17

If \tilde{T} is an expanded signature set for x_j , and $\tilde{A}_{\tilde{T}}$ is the corresponding column output by Algorithm 1, then with high probability $\|\tilde{A}_{\tilde{T}} - A_j\|_\infty \leq O(\rho(\Lambda^3 \log n / \sigma^2)^2 \sqrt{\Lambda \log n}) = o(\sigma)$.

Once we have all the entries which are $> \sigma/2$ in magnitude, we can use the ‘refinement’ trick of Lemma 13 to conclude that we can recover the entries.

LEMMA 18

When the number of samples is at least $n^{4C+3}m$, the matrices A and \hat{A} are entry-wise $n^{-2C}m^{-1/2}$ close. Further, the two dictionaries are n^{-C} -equivalent.

The first part of the proof (showing entry-wise closeness) is very similar to Lemma 6. In order to show n^{-C} equivalent, notice when the entries are very close this just follows from Bernstein’s inequality, with variance bounded by $n^{-4C}m^{-1} \cdot m$. In Section 3 we do not just use this bound, because we want to be able to also handle the case when the entrywise error is only inverse polylog (for Assumption 2).

References

- [AAN13] Alekh Agarwal, Animashree Anandkumar, and Praneeth Netrapalli. Exact recovery of sparsely used overcomplete dictionaries. *CoRR*, abs/1309.1952, 2013.
- [ABGM13] Sanjeev Arora, Aditya Bhaskara, Rong Ge, and Tengyu Ma. Provable bounds for learning some deep representations. *CoRR*, abs/1310.6343, 2013.
- [AEB05] Michal Aharon, Michael Elad, and Alfred M Bruckstein. K-svd and its non-negative variant for dictionary design. In *Optics & Photonics 2005*, pages 591411–591411. International Society for Optics and Photonics, 2005.
- [AEB06] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *Signal Processing, IEEE Transactions on*, 54(11):4311–4322, 2006.
- [AEP06] Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In *NIPS*, pages 41–48, 2006.
- [AGM13] Sanjeev Arora, Rong Ge, and Ankur Moitra. New algorithms for learning incoherent and overcomplete dictionaries. *ArXiv*, 1308.6273, 2013.

- [Aha06] Michal Aharon. *Overcomplete Dictionaries for Sparse Representation of Signals*. PhD thesis, Technion - Israel Institute of Technology, 2006.
- [BC⁺07] Y-lan Boureau, Yann L Cun, et al. Sparse feature learning for deep belief networks. In *Advances in neural information processing systems*, pages 1185–1192, 2007.
- [Ben62] George Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):pp. 33–45, 1962.
- [Ber27] S. Bernstein. *Theory of Probability*, 1927.
- [BGI⁺08] R. Berinde, A.C. Gilbert, P. Indyk, H. Karloff, and M.J. Strauss. Combining geometry and combinatorics: a unified approach to sparse signal recovery. In *46th Annual Allerton Conference on Communication, Control, and Computing*, pages 798–805, 2008.
- [CRT06] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [Das99] Sanjoy Dasgupta. Learning mixtures of gaussians. In *FOCS*, pages 634–644. IEEE Computer Society, 1999.
- [DH01] David L Donoho and Xiaoming Huo. Uncertainty principles and ideal atomic decomposition. *Information Theory, IEEE Transactions on*, 47(7):2845–2862, 2001.
- [DMA97] Geoff Davis, Stephane Mallat, and Marco Avellaneda. Adaptive greedy approximations. *Constructive approximation*, 13(1):57–98, 1997.
- [EA06] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *Image Processing, IEEE Transactions on*, 15(12):3736–3745, 2006.
- [EAHH99] Kjersti Engan, Sven Ole Aase, and J Hakon Husoy. Method of optimal directions for frame design. In *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*, volume 5, pages 2443–2446. IEEE, 1999.
- [Hoy02] Patrik O Hoyer. Non-negative sparse coding. In *Neural Networks for Signal Processing, 2002. Proceedings of the 2002 12th IEEE Workshop on*, pages 557–565. IEEE, 2002.
- [Ind08] Piotr Indyk. Explicit constructions for compressed sensing of sparse signals. In Shang-Hua Teng, editor, *SODA*, pages 30–33. SIAM, 2008.
- [JXHC09] Sina Jafarpour, Weiyu Xu, Babak Hassibi, and A. Robert Calderbank. Efficient and robust compressed sensing using optimized expander graphs. *IEEE Transactions on Information Theory*, 55(9):4299–4308, 2009.

- [LS99] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- [LS00] Michael S Lewicki and Terrence J Sejnowski. Learning overcomplete representations. *Neural computation*, 12(2):337–365, 2000.
- [MLB⁺08] Julien Mairal, Marius Leordeanu, Francis Bach, Martial Hebert, and Jean Ponce. Discriminative sparse image models for class-specific edge detection and image interpretation. In *Computer Vision–ECCV 2008*, pages 43–56. Springer, 2008.
- [OF97] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.
- [SWW12] Daniel A. Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. *Journal of Machine Learning Research - Proceedings Track*, 23:37.1–37.18, 2012.
- [YWHM08] Jianchao Yang, John Wright, Thomas Huang, and Yi Ma. Image super-resolution as sparse representation of raw image patches. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

A Full Proofs

In this section we give the omitted proofs.

A.1 Proof of Lemma 14

PROOF: Let us focus on the i th row of $A - \hat{A}$ and denote it by w . Then we have $\|w\|_1 \leq \|A\|_1 + \|\hat{A}\|_1 \leq O(1/\rho)$. Now consider the random variable $Z = \sum_j w_j x_j$, where x_j are i.i.d. Bernoulli r.v.s with probability ρ of being 1. Then by Bernstein’s inequality (Theorem 23), we have

$$\Pr[Z - \mathbb{E} Z > \epsilon] \leq e^{-\frac{\epsilon^2}{\rho\sigma + \rho \sum_j w_j^2}}.$$

Since $|w_j| < \delta$ for all j , we can bound the variance as $\rho \cdot \sum_j w_j^2 \leq \delta\rho \cdot \sum_j |w_j| \leq 2\delta$.

Thus setting $t = (4\delta \log n)^{1/2}$ (notice that this is the t in Bernstein’s inequality, not the same as the size of signature sets), we obtain an upper bound of $1/\text{poly}(n)$ on the probability. \square

A.2 Missing Lemmas and Proofs of Section 4

LEMMA 19 (GENERAL VERSION OF LEMMA 3)

Suppose T of size t is a general signature set for x_j with $t = \omega(\sqrt{\log n})$. Let E_1 be the event that $x_j = 1$ and E_2 be the event that $\beta_T \geq \mathbb{E}[\beta_T] + 0.9\sigma t$ if $\beta_{j,T} \geq \sigma t$, and the event $\beta_T \leq \mathbb{E}[\beta_T] - 0.9\sigma t$ if $\beta_{j,T} \leq -\sigma t$. Then for large constant C (depending on Δ)

1. $\Pr[E_1] + n^{-2C} \geq \Pr[E_2] \geq \Pr[E_1] - n^{-2C}$.
2. $\Pr[E_2|E_1] \geq 1 - n^{-2C}$, and $\Pr[E_2|E_1^c] \leq n^{-2C}$.
3. $\Pr[E_1|E_2] \geq 1 - n^{-C}$.

PROOF: It is a straightforward modification of the proof of Lemma 3. First of all, $|\mathbb{E}[\beta_{j,T}x_j]| = o(\sigma t)$, and thus mean of $\sum_{k \neq j} \beta_{k,T}x_k$ only differs from that of β_T by at most $o(\sigma t)$. Secondly, Bernstein inequality requires the largest coefficients and the total variance be bounded, which correspond to exactly property (b) and (c) of a general signature set. The rest of the proof follows as in that for Lemma 3. \square

LEMMA 20 (GENERAL VERSION OF LEMMA 8)

Let \tilde{T}^* be the set with largest general empirical bias $\hat{B}_{\tilde{T}^*}$ among all the expanded sets \tilde{T} . The set \tilde{T}^* is an expanded signature set for some x_j .

PROOF: We first prove an analog of Claim 9. Let $W = \sigma^4 d / 2\Delta \Lambda^3 \log n$. Let's redefine $K_{large} := \{k \in [m] : |\{i \in \tilde{T} : |A_k^{(i)}| \geq \tau\}| \geq W\}$ be the subset of nodes in $[m]$ which connect to at least W nodes in \tilde{T} in the subgraph G_{τ} . Note that this implies that if $k \notin K_{large}$, then $|\beta_{k,\tilde{T}}| \leq d\tau + W\Lambda \leq d\sigma^4 / \Delta \Lambda^2 \log n$. Let $Q_k = \{i \in \tilde{T} : |A_k^{(i)}| \geq \tau\}$. By definition, we have for $k \in K_{large}$, $|Q_k| \geq W$. Then similarly as in the proof of Claim 9, using the fact that $|Q_k \cap Q_{k'}| \leq \kappa$, and inclusion-exclusion, we have that $|K_{large}| \leq O(\Delta \Lambda^3 \log n / \sigma^4)$.

Then we prove an analog of Claim 10. Let $\beta_{small,\tilde{T}}$ and $\beta_{large,\tilde{T}}$ be defined as in the proof of Claim 10 (with the new definition of K_{large}). By Lemma 15, the variance of $\beta_{small,\tilde{T}}$ is bounded by $2dW + 2d^2\gamma \leq 2d^2\sigma^4 / \Delta \Lambda^2 \log n$. Therefore by Bernstein's inequality we have that for sufficiently large Δ , with probability at least $1 - n^{-2}$ over the choice of x , $|\beta_{small,\tilde{T}} - \mathbb{E}[\beta_{small,\tilde{T}}]| \leq 0.05d\sigma^2 / \Lambda$. It follows from the same argument of Claim 10 that with high probability over the choice of N samples, $|\hat{B}_T - \max_k \beta_{k,\tilde{T}}| \leq 0.1d\sigma^2 / \Lambda$ holds when $\max_k \beta_{k,\tilde{T}} \geq 0.5d\sigma$.

We apply almost the same argument as in the proof of Lemma 8. By Lemma 17 we know that our algorithm must produce an expanded signature set of size d with bias at least $0.8\sigma d$, and thus the set \tilde{T}^* with largest bias must has a large coefficient j with $\beta_{j,\tilde{T}^*} \geq 0.7\sigma d$. If there is some other k such that $\beta_{k,\tilde{T}^*} \geq 0.3\sigma d$, then $|Q_k| \geq 0.3\sigma d / \Lambda$ and therefore we could remove those elements in $\tilde{T}^* - Q_j$, which has size larger than $0.3\sigma d / \Lambda - \kappa$ by Assumption G2. Then by adding some other elements which are in the neighborhood of j in G_{σ} into the set Q_j we get a set with bias larger than \tilde{T}^* , which contradicts our assumption that there exists k with $\beta_{k,\tilde{T}^*} \geq 0.3\sigma d$. Hence \tilde{T}^* is indeed an expanded signature set and the proof is complete.

\square

B Probability Inequalities

LEMMA 21

Suppose X is a bounded random variable in a normed vector space with $\|X\| \leq M$. If event E happens with probability $1 - \delta$ for some $\delta < 1$, then $\|\mathbb{E}[X|E] - \mathbb{E}[X]\| \leq 2\delta M$

PROOF: We have $\mathbb{E}[X] = \mathbb{E}[X|E] \Pr[E] + \mathbb{E}[X|E^c] \Pr[E^c] = \mathbb{E}[X|E] + (\mathbb{E}[X|E^c] - \mathbb{E}[X|E]) \Pr[E^c]$, and therefore $\|\mathbb{E}[X|E] - \mathbb{E}[X]\| \leq 2\delta M$. \square

LEMMA 22

Suppose X is a bounded random variable in a normed vector space with $\|X\| \leq M$. If events E_1 and E_2 have small symmetrical differences in the sense that $\Pr[E_1|E_2] \leq \delta$ and $\Pr[E_2|E_1] \leq \delta$. Then $\|\mathbb{E}[X|E_1] - \mathbb{E}[X|E_2]\| \leq 4\delta M$.

PROOF: Let $Y = X|E_2$, by Lemma 21, we have $\|\mathbb{E}[Y|E_1] - \mathbb{E}[Y]\| \leq 2\delta M$, that is, $\|\mathbb{E}[X|E_1E_2] - \mathbb{E}[X|E_2]\| \leq 2\delta M$. Similarly $\|\mathbb{E}[X|E_1E_2] - \mathbb{E}[X|E_1]\| \leq 2\delta M$, and hence $\|\mathbb{E}[X|E_1] - \mathbb{E}[X|E_2]\| \leq 4\delta M$. \square

THEOREM 23 (BERNSTEIN INEQUALITY [BER27] CF. [BEN62])

Let x_1, \dots, x_n be independent variables with finite variance $\sigma_i^2 = \mathbb{V}[x_i]$ and bounded by M so that $|x_i - \mathbb{E}[x_i]| \leq M$. Let $\sigma^2 = \sum_i \sigma_i^2$. Then we have

$$\Pr \left[\left| \sum_{i=1}^n x_i - \mathbb{E} \left[\sum_{i=1}^n x_i \right] \right| > t \right] \leq 2 \exp \left(- \frac{t^2}{2\sigma^2 + \frac{2}{3}Mt} \right)$$